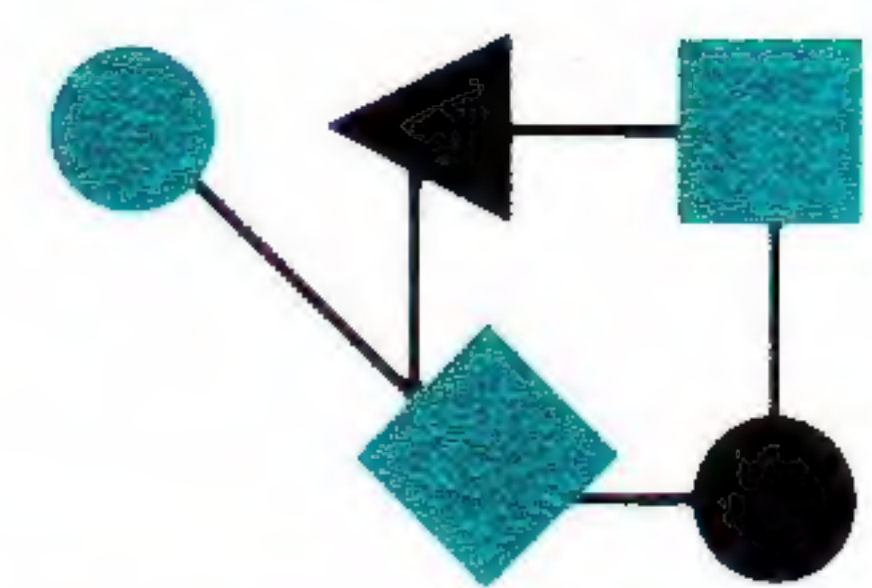


CONNEXIONS



The Interoperability Report

September 1988

Volume 2, No. 9

*ConneXions -
The Interoperability Report
tracks current and emerging
standards and technologies
within the computer and
communications industry.*

In this issue:

The HP Internet.....	2
Amateur Packet Radio and TCP/IP.....	8

From the Editor

We have often spoken about how TCP/IP provides true interoperability in multi-vendor environments. TCP/IP's strength lies in its ability to run over every imaginable type of media, from slow telephone circuits to Ethernets and satellite networks. By connecting such diverse technologies together by means of IP routers (or "gateways" in the Arpanet jargon) one can construct arbitrarily large *internets*. In this issue we will look at two very different examples of such internets. The first is a "User Success Story:" Hewlett-Packard's Internet is the largest TCP/IP-based private corporate network in the world with some 6500 hosts. Norm Kincl and Robert Michaels of HP give an overview on the design and use of this internet.

In the second article, Phil Karn describes how radio amateurs are using TCP/IP to build ad-hoc internets. The radio "ether" is much more hostile than your average backbone network and special methods must be employed to ensure proper operation of TCP/IP over what the amateurs refer to as the "Worst Wire." The lessons learned from these experiments have become valuable to the TCP/IP community as a whole since many of the algorithms developed have applicability outside the amateur packet radio field.

Advanced Computing Environments' headquarters in Mountain View, California is a busy place these days as we enter the final 3 week countdown for *INTEROP 88: The 3rd TCP/IP Interoperability Conference & Exhibition* which will be held September 26 - 30. As of this writing we have over 50 vendors signed up for the exhibition which promises to be an exciting event. The Netman Working Group of the Internet Engineering Task Force (IETF) is also busy putting the final touches on their demonstration of a network management system based on the ISO CMIS/CMIP standard to be shown at INTEROP 88. A number of people have contacted us to schedule Birds Of a Feather (BOF) sessions at the conference. If you have a topic which you think is suitable for a BOF, please give us a call at 415-941-3399 and we will take care of the arrangements.

As scheduled, by the beginning of July 1988 the new NSFNET backbone began operational services. A new era of national T1 networking got started by this milestone event. A few weeks of successful operation of the new network allowed for the complete phaseout of the old, 56Kbps based, NSFNET backbone. We expect detailed reports of the new NSFNET backbone in subsequent issues of *ConneXions*.

ConneXions is published by Advanced Computing Environments, 480 San Antonio Road, Suite 100, Mountain View, CA 94040, USA. Phone: 415-941-3399.

© 1988
Advanced Computing Environments.
Quotation with attribution encouraged.

ConneXions-The Interoperability Report
and the *ConneXions* masthead are
trademarks of Advanced Computing
Environments.

ISSN 0894-5926

Inside Autonomous System 71 - The HP Internet

by Norman Kincl and Robert Michaels,
Hewlett-Packard Laboratories

Background

In August of 1985 Hewlett-Packard decided to solve a communication problem in our R&D community. At the time, all data communication was done with *uucp* using 1200 baud modems or X.25 at up to 9600 baud, often much slower. These methods were expensive, slow, and inadequate. Monthly communications charges at some sites were several thousand dollars for next-day *uucp* service.

What we needed was a high-bandwidth, interactive network connecting the R&D facilities in the company. The initial pilot plan called for connecting HP's R&D sites in Palo Alto, CA with Cupertino, CA, Ft. Collins, CO and Corvallis, OR. The eventual goal was to connect facilities worldwide.

Design

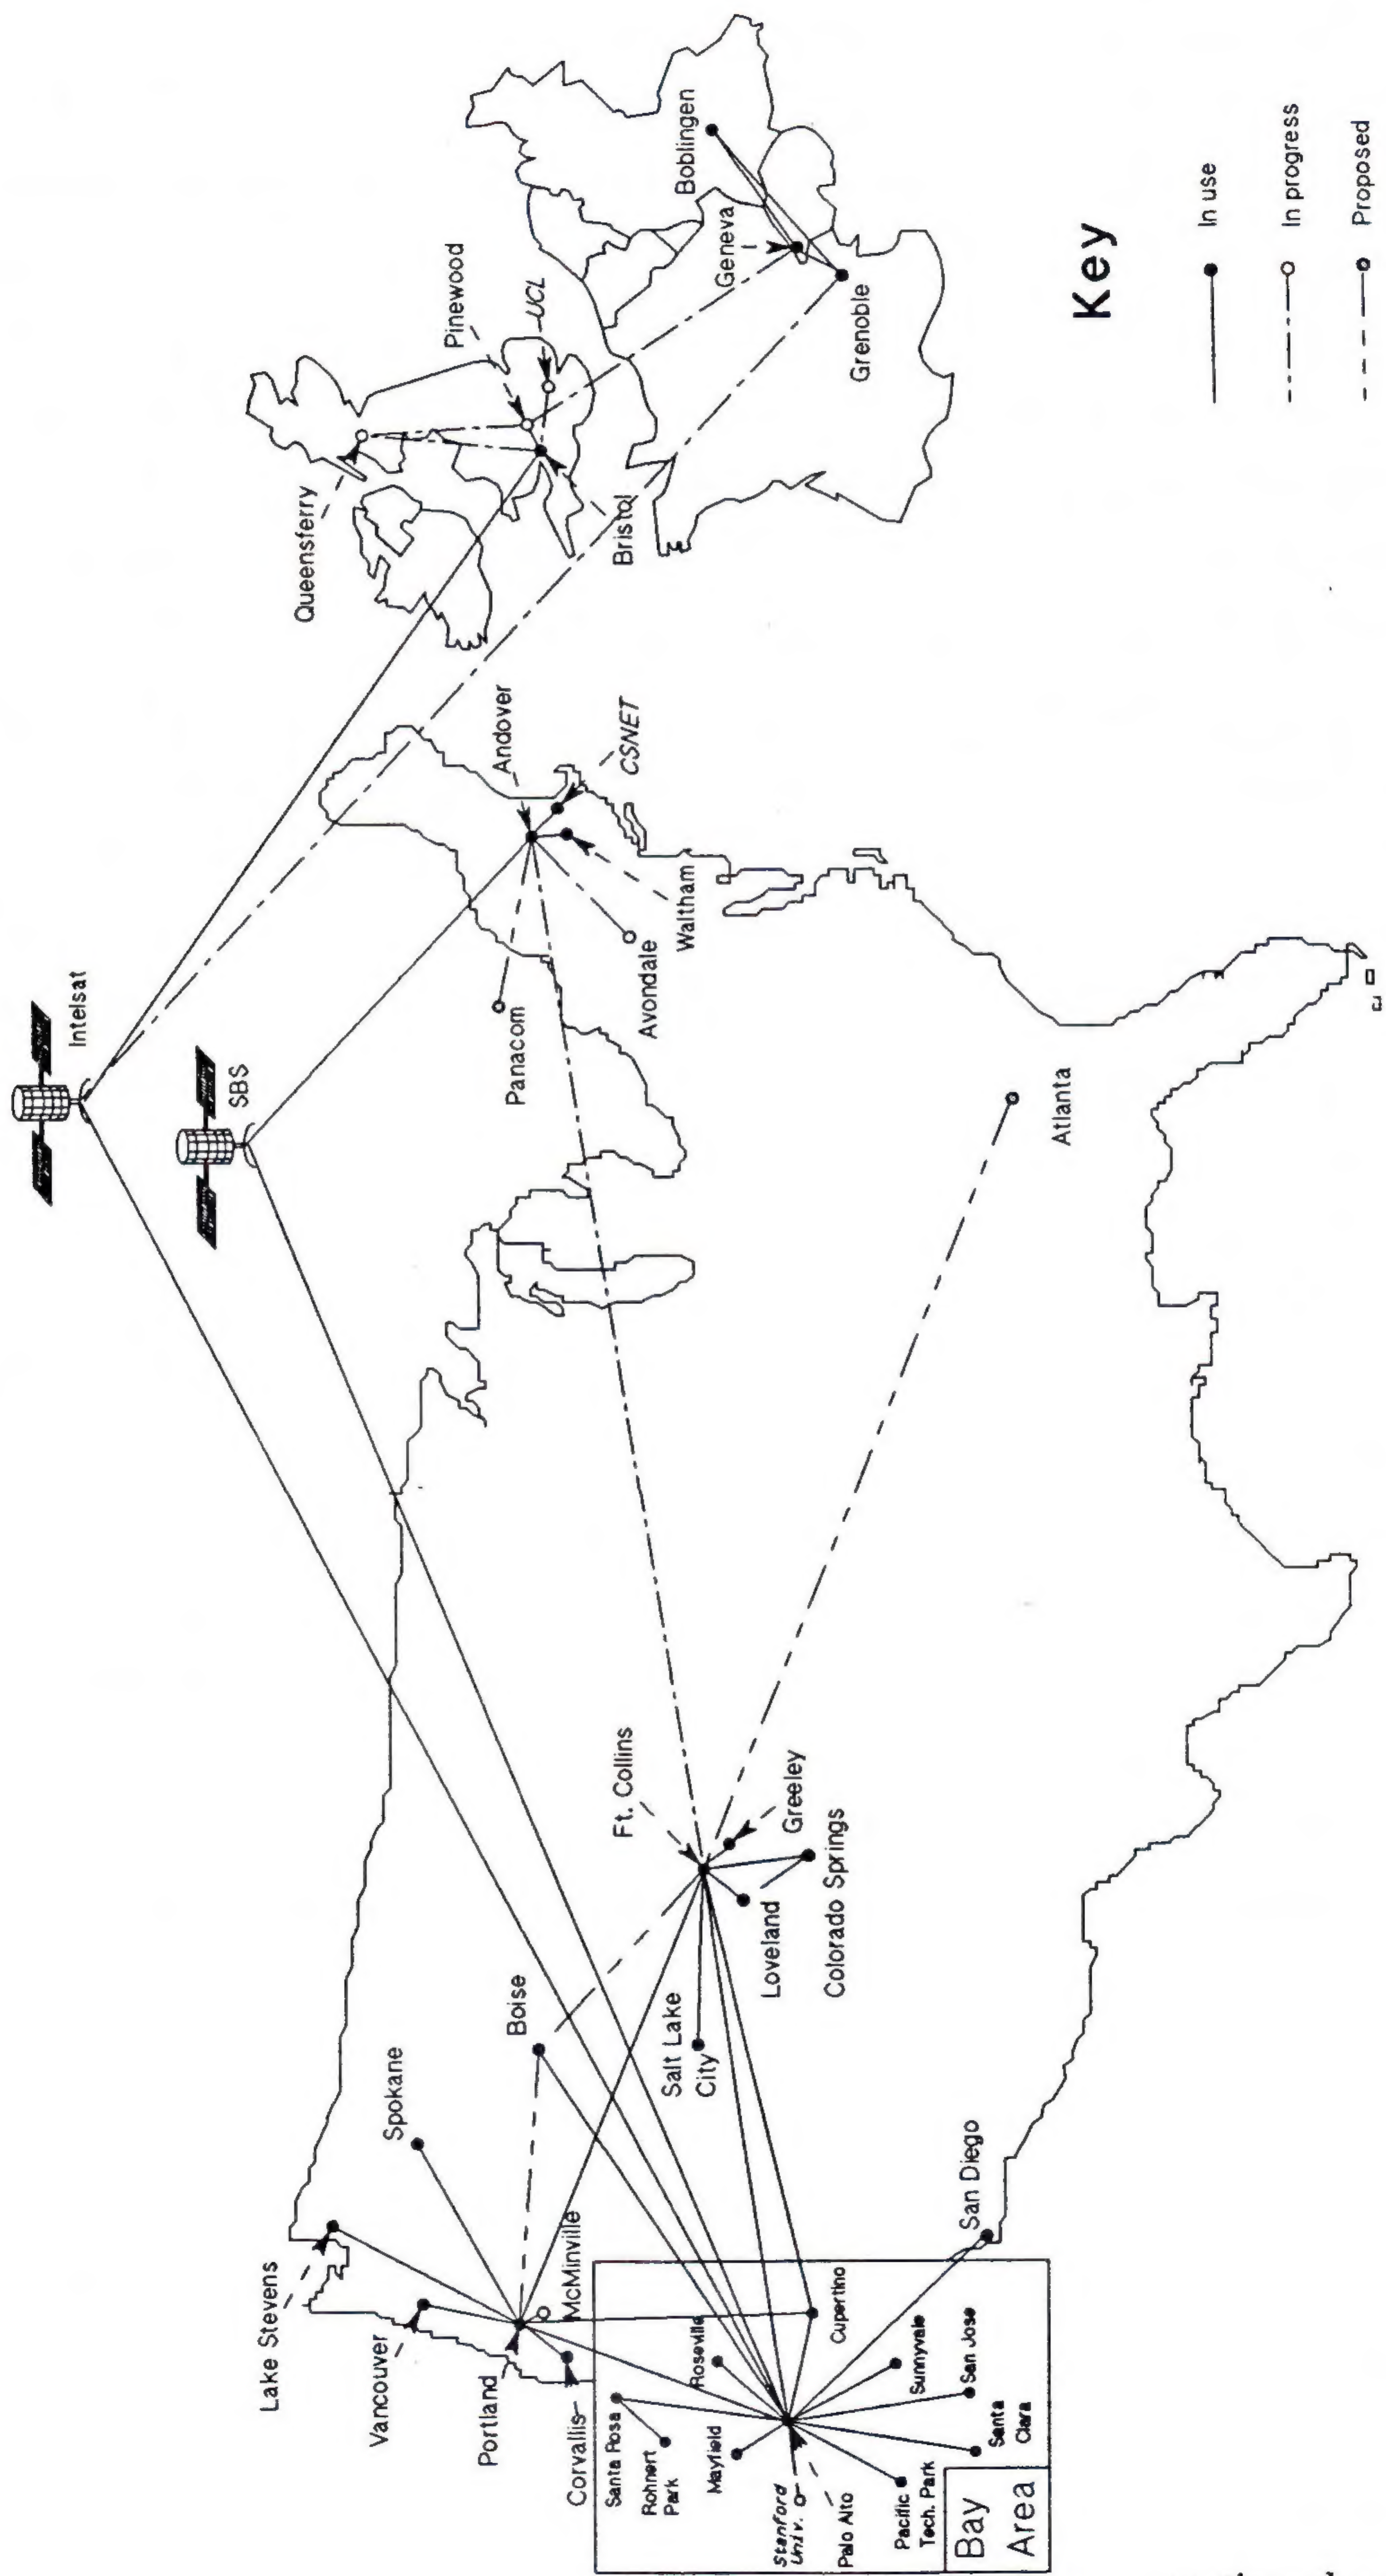
Since Local Area Networks (LANs) were proliferating rapidly throughout HP's engineering community the logical solution was to interconnect these networks. We examined two possibilities - the use of bridges such as Vitalink's TransLAN and the use of level 3 routers. After a careful evaluation of the benefits of bridges and routers, we decided to build the Wide Area Network (WAN) using IP based routers. [Bosack and Hedrick] provide a good comparison of bridges and routers.

We were confident that a network built on level 3 routers would work well. The best example we had to look at was the Arpanet. IP gateways had been in use on the Arpanet for several years. ("Gateway" is the Arpanet community term for the ISO term "Level 3 Router." We will use gateway and router interchangeably in this article.) They allow a combination of highly interactive networking and access from every computer in a geographically widely distributed environment.

Providing the HP engineering community with an IP based internet also has significant strategic advantages. It ensures that our TCP/IP products can operate well in a complex network. It also is providing us with significant experience with heterogeneous networking protocols. This experience will prove valuable as we move to the ISO protocols in the next several years.

Growth

We evaluated IP gateway vendors in the spring of 1986. We began building the network that summer, bringing up the first five sites in July 1986. By January 1987, the HP Internet reached twenty divisions and over 1,000 hosts. In August of 1987 we connected the HP Internet to CSNET with a high speed satellite connection, and through CSNET to the thousands of hosts on the DARPA Internet. By October 1987, forty divisions and over 4,000 hosts within HP had access to the network. We now have over 6,500 hosts on 210 subnets connected by about 120 gateways. As far as we know, HP Internet is the largest TCP/IP network operated by a single organization. The map on the next page shows the current extent of the network.----->



Key

- In use
- - -○- - - In progress
- - -●- - - Proposed

 Internet May 1988

continued on next page

Inside Autonomous System 71 (*continued*)

Since the network currently serves HP's engineering community, the predominant hosts on the network are HP9000 and HP Vectra computers. Additionally there are HP1000s and HP3000s, as well as systems from many other vendors.

The explosive growth shows just how desperately the HP R&D community needed the network. It also showed us that our initial decisions on how to build the network were correct.

Construction details

The HP Internet can be divided into two basic components: the *gateways* and the *links* that connect them. The gateways are the key devices which make all the routing decisions and move the traffic from one link to another. The HP Internet uses a wide variety of links to interconnect the gateways.

Gateway selection

Using a single vendor's gateway throughout the network eases management problems and allows for a straightforward plan for expansion. Initial plans for the network called for linking the sites together using microVAXen. The VAX workstation would run standard Berkeley Software Distribution (BSD) software, which has gateway functionality. However, this code has a history of problems and, given the high cost of these workstations, we needed a more reliable, lower cost alternative.

A dedicated gateway solution appeared to be the best alternative. After looking at two vendor's products, cisco Systems of Menlo Park, CA was chosen. The system had many features that were desirable:

- *RFC 1009*: Though RFC 1009 was not yet written at the time of our selection, we identified most of the requirements that are now in that document.
- *Routing*: cisco uses its own proprietary routing protocol to automatically spread information about the various subnets of the HP Internet to all the gateways. Each gateway gains information about the structure of the network by listening to its neighboring gateways. Such things as bandwidth, delay and reliability are included in the metrics passed for each route. This allows a gateway to make intelligent decisions about how to route packets. If two or more paths exist between two parts of the network, the gateways will dynamically load balance between paths. cisco gateways will also speak other routing protocols such as RIP [RFC 1058], HELLO [RFC 891], and EGP [RFC 904] if required.
- *Management*: The gateways can be fully controlled and managed using the standard Telnet [RFC 854] protocol. Each gateway boots through the network from a central server or servers (typically a workstation). The servers can be anywhere on the network. The standard protocols TFTP [RFC 783] and BOOTP [RFC 951] are used. Overall network management is a goal for HP Internet. We are anxiously awaiting the results of the efforts to develop standards for TCP/IP network management, see [RFC 1052].

- *Futures:* Because of HP's commitment to move to ISO standards, we needed a vendor committed to supporting the ISO End System-to-Intermediate System (ES-IS) protocol. cisco Systems has since announced this as part of their product.

Links and interfaces

The HP Internet is built on a variety of media or physical links. In addition to LAN (Ethernet and IEEE 802.3) the network also uses 56 Kbit and 1.544 Mbit (T1) land line serial links, IP on X.25, broadband, and serial satellite links.

To link divisions which are physically far apart we made use of HP's internal T carrier network whenever possible. In a couple of cases satellite links were necessary. The T carrier system breaks up a T1 link into twenty-four 56 Kbit channels to carry voice or data traffic.

Satellite

Two of the longest paths within the network use satellite connections. One link is from California to Massachusetts and the other is from California to HP's research facility in Bristol, England. In the case of the domestic connection, HP already owned earth stations at both sites to support our video conferencing facility. For roughly one fifth the monthly cost of a transcontinental 56 Kbit Digital Data Service (DDS) circuit, HP could outfit and operate the 56 Kbit satellite circuit. The savings justifies the inconvenience of the 800 millisecond round-trip satellite delay. The link to Bristol is an end-to-end 64 Kbit per second circuit provided by ITT World Communications.

X.25

In certain cases we have found it advantageous to use HP's large, internal X.25 network. Two gateways can establish an X.25 circuit between themselves and use it to switch IP packets. This is especially useful in the following situations:

- Provide a redundant link as a backup to the dedicated IP circuits.
- Provide for connections to sites with low IP traffic.
- Provide connections to sites where dedicated circuits are expensive. We plan to initially connect the Pacific rim sites in this manner.
- Provide for temporary connections. We have done this both for demos and also to connect sites while awaiting installation of the leased line.

For campus-wide environments most sites use baseband Ethernet or 802.3 installations. In Palo Alto, a mid-split broadband system is used to carry traffic between buildings.

Security

The cisco gateway provides the capability to limit certain source and/or destination addressed IP packets from being sent out a particular interface. Additionally, one can filter on different protocol types along with certain port numbers within this protocol. This gives a lot of flexibility for providing security. Using these features, we have defined "open" and "closed" subnets - those that can and those that cannot be accessed from networks outside HP. We gain additional security by fully encrypting the satellite links.

continued on next page

Inside Autonomous System 71 *(continued)*

Use The HP Internet has provided for significant productivity increase among our engineers. We have documented cases of productivity increases [Cook] from all types of engineers. A few examples are provided here.

The network has been used to reduce travel: "Last year I flew from Cupertino to Boise to do some performance measurement... Shortly after they got on our internet I was able to perform similar tests in Boise while sitting at my desk in California."

It is used to distribute software: "The HP Internet is very nice for Alpha site distribution. I used the internet to distribute Alpha versions of HP 9000/300 HP-UX Pascal 6.0. For the internet sites, I just made a *cpio* archive of the Pascal bits which were then available via *ftp* (later, I made Pascal 300 installable) [Jin]. For the sites not connected to the internet, I had to distribute... tapes which required much more time on my part and the time of HP clerical personnel. The distribution was much slower, too."

We share source code across country: "All of the shared HP-UX commands source..., for which the master copy is maintained at Information Software Operation (ISO) [in Cupertino], is synchronized with our copy in Fort Collins via the internet. Each day we exchange any RCS files which don't match at the two sites. Every time we do a revision here, it is immediately sent to ISO."

The interactive nature of HP Internet can even help with such things as IC design: "When the part came out of fab, it was time to try out the test program, using the vectors that they had sent me over the net. Over the next 8 days, the test program and vectors were debugged to the point where we were able to ascertain that both the chip and the test program were working. We sent them proto parts, which they plugged into their prototype tape drives, which worked first time! The only reason that the debug took only 8 days was because the [Greeley] engineers were able to sit at their desks in Colorado and *watch in real time* the output of the tester here in Corvallis...

Now, here's the clincher: if the wafers had come out of fab one month earlier, then the [Greeley] engineers would have had to fly out here to help me debug. They would also have had to bring as much of their online design data as possible, and we would have had to find a machine here to put it on. In summary, because of the internet... we were able to save many man-days worth of work and plane fares on this project.

Management (4 levels up!) have asked me to write an article... describing in more detail what I just told you. If you're looking for people convinced that Internet is an asset, this is the place."

Futures The next step in HP Internet has been connecting the various sites in Europe. This is currently in progress. By the time this article is published, a second transatlantic satellite link (Palo Alto to Grenoble) should be installed. New higher-speed links connecting Geneva and the United Kingdom will also be in place shortly. We expect that the satellite connections will be replaced by transatlantic fiber service once TAT-8 is operational.

After connecting Europe, we will start work on connecting Japan, Singapore and Australia. Initially, connections will be via our X.25 network. As utilization increases and additional bandwidth is installed across the Pacific, we will move to dedicated links.

A long term goal is to provide complete redundancy within HP Internet. We are planning on several different approaches for redundancy. Some of it will be provided by additional use of satellite links. Other redundancy will be implemented as part of the X.25 network. Finally, lower bandwidth links may be installed between divisions which currently are not directly connected. For example, a direct connection between Colorado and Oregon would mean that traffic no longer needs to go through California.

References

[Bosack and Hedrick] Bosack, L. and Hedrick, C., "Problems in Large LANs," IEEE Network, Volume 2, No. 1, Jan 1988.

[Cook] Cook, T. and Underwood, W., "Experience with the HP Internet," European Software Engineering Productivity Conference, Hewlett-Packard, November 1987. This is an HP internal document.

[Jin] Jin, T., Kincl, N., Thompson, B., and Untulis, C., "Network Based Software Distribution," UniForum Conference, /usr/group, Feb 1988.

[Michaels] Michaels, R., Kincl, N. and Underwood, W., "Designing and Installing an International TCP/IP Network," UniForum Conference, /usr/group, Feb 1988.

[Kincl] Kincl, N., Jin, T. and Michaels, R., "Managing a Distributed Environment," Proc. Software Engineering Productivity Conference, Hewlett-Packard, May 1987. This is an HP internal document.

[RFC 783] Sollins, K. R., "The TFTP Protocol," June 1981.

[RFC 854] Postel, J. and Reynolds, J. "Telnet Protocol Specification," May 1983.

[RFC 891] Mills, D. L., "DCN Local-Network Protocols," Dec 1983.

[RFC 904] Mills, D. L., "Exterior Gateway Protocol Formal Specification," April 1984.

[RFC 951] Croft, W. and Gilmore, J., "Bootstrap Protocol," Sept 1985.

[RFC 959] Postel, J. and Reynolds, J., "File Transfer Protocol (FTP)," Dec 1985.

[RFC 1009] Braden, R. and Postel, J., "Requirements for Internet Gateways," June 1987.

[RFC 1052] Cerf, V., "IAB Recommendations for the Development of Internet Network Management Standards," April 1988.

[RFC 1058] Hedrick C. L., "Routing Information Protocol," June 1988.

Amateur Packet Radio and TCP/IP

by Phil Karn, KA9Q

Introduction

A new and rapidly growing activity within amateur (ham) radio is *packet radio*. From a small-scale experimental beginning in the early 1980s, an estimated 30,000 amateurs around the world now own equipment capable of reliably sending data over the air. [Karn]

Of course, hams did not invent packet radio; the Aloha experiments in Hawaii predated the earliest amateur networks by almost a decade. However, amateur packet radio is noteworthy for the following reasons:

Frequency coordination

[1] Amateur radio operation requires a government license, granted only to applicants passing tests in Morse code, radio and electronics theory, and rules and regulations. Once licensed, however, a ham has easy access to a wide range of radio frequency bands, from just above the AM broadcast band well up into the microwave region. Within broad policy and technical limits established by government rules, amateurs have considerable freedom to experiment. Frequency coordination within the amateur bands is handled largely by the amateurs themselves, in strong contrast to the various non-amateur, non-military packet radio experiments that often appear inhibited by regulatory red tape.

Operating modes

[2] Amateur radio includes many operating modes and propagation techniques. Packet radio coexists with more traditional modes such as analog voice (single-sideband and FM), Morse code, facsimile, radioteletype and television. Signals are propagated directly or by the ionosphere, orbiting satellite or local hilltop repeater, and even the moon is occasionally used as a passive reflector! These modes support a broad range of applications, from casual socializing to technical experimentation for its own sake, to public service in time of emergency. Business use (profit or non-profit) is prohibited, however, and amateurs may not charge or accept payment for the communication services they provide.

Hackers

[3] Similarly, the amateurs themselves are a varied lot. Some are communications professionals but many are not. Amateur radio has a strong tradition of close-knit volunteer cooperation, strengthened by the knowledge that no one will profit financially from the efforts of others, at least not directly. Amateurs therefore must be strongly motivated by an interest in communications technology for its own sake. There is a strong parallel between many amateur packet radio experimenters and the early university computer science researchers described in Steven Levy's book "Hackers" [Levy]; indeed, many *are* computer hackers in their own right, because many of the same skills are required. (Here I use the original meaning of the word: one who loves programming for its own sake, not one who gains unauthorized access to computers.)

Low cost

[4] Amateur radio equipment is extremely inexpensive by commercial and military standards because it must be affordable by ordinary individuals to whom communications and computer networking is a hobby, not a profession. As a new mode, amateur packet radio has emphasized using equipment the average amateur is already likely to have: a voice transceiver, and a terminal or small personal computer. The initial emphasis has been on low cost rather than performance and efficiency.

The state of Amateur Packet Radio

The packet radio techniques now in widespread amateur use reflect this low cost philosophy. Almost all VHF/UHF amateur packet radio currently uses Bell 202-compatible modems, as the 1200/2200 Hz frequency-shift-keyed (FSK) tones they produce are easily fed into standard FM voiceband radios. On the "short wave bands," below 30 Mhz, a crowded spectrum calls for more efficiency: standard FSK radioteletype modems with 170 Hz shift are used with single side-band (SSB) radios. A major development is underway that will apply digital signal processing (DSP) to amateur radio, with a goal being the creation of more efficient low speed HF modems. [McGwier]

When amateur packet radio began, there were few users and PCs were much rarer than "dumb terminals." Slow modems could support the simple keyboard-to-keyboard "chatting" that was the usual operating style. Increased activity along with the need to transfer electronic mail and files has spurred development of higher speed modems. One available unit uses Minimum Shift Keying and operates at 56Kbps. Such modems are restricted to the higher frequency bands due to the bandwidth required. [Heatherington]

TNC

The packet radio modem is generally incorporated into a special-purpose box called, for historical reasons, a *Terminal Node Controller* (TNC). TNCs also contain single-board computers with the necessary firmware to execute the packet protocols. One port on the TNC connects to the radio's audio I/O and push-to-talk leads, while an RS-232 connection is provided for the host terminal or computer. A TNC is generally about the same size as a small telephone modem, although one new Japanese model is not much larger than a typical RS-232 "null modem." Lately HDLC adapter/modem cards for standard personal computers have appeared, eliminating the need for separate TNCs altogether.

Radio is *not* wire

It must be emphasized that the packet radio channel is radically different from the far more benign wire or fiber transmission path, and this heavily influences the design of a network and its lower-level protocols. Although a collection of packet radio stations operating on a single, shared frequency channel bears a strong resemblance to multiple-access local area networks such as Ethernet (which, conversely, has been described as "packet radio on a cable") there are some important differences:

[1] Not every station can hear the transmissions of every other station. This violates one of the fundamental assumptions of the Carrier Sense Multiple Access technique; the ability to defer transmission when the channel is busy. Known as the "Hidden Terminal" problem, this can severely degrade network efficiency.

[2] It is impossible for a station to monitor the channel while transmitting because of the enormous difference between received and transmitted signal levels (150 dB is not unusual). This rules out Collision Detection a la Ethernet.

[3] Depending on the modem hardware and RF path, the bit error rate of the channel may range from somewhat worse to much worse than a wire or fiber path.

[4] Due to bandwidth and power limitations, the signaling speed of the packet channel is usually much lower than a local area network.

Amateur Packet Radio and TCP/IP (*continued*)

Attacking these problems is now a major activity within amateur packet radio. For example, it is possible to eliminate hidden terminals and to detect collisions if a full duplex RF repeater is used so that the stations can transmit and receive on separate frequencies simultaneously. Or the network can be engineered such that each transmitter is assigned its own frequency, with each node listening to its neighbors with separate receivers.

Link protocols

Since 1982, the standard amateur link level protocol has been "AX.25" [ARRL82]. It must be stressed that this name is somewhat of a misnomer; unlike BX.25 ("Bell System X.25"), AX.25 it is *not* compatible with CCITT X.25. AX.25 was custom-designed for amateur radio, although it does incorporate many elements of X.25's link layer, LAPB (Link Access Procedures Balanced).

As in X.25, AX.25 packets are sent synchronously in HDLC frames. (At 1200 baud, the lack of start and stop bits gives a welcome, though modest, performance improvement!) Since AX.25 is used on a shared-access channel, half-duplex operation is the norm. A datagram-style address header is included in the beginning of each frame. As a minimum, the header contains the FCC-assigned callsigns of the destination and source stations in ASCII; this has the side benefit of satisfying FCC identification rules. Beyond the source field may be up to eight "digipeater" addresses. A digipeater is an intermediate station capable of receiving, storing and retransmitting a packet, usually on the same frequency. Strict source routing is used; the sender must specify the entire digipeater string.

Beyond the address field is the LAPB control field and a protocol ID (PID) byte. The PID corresponds to the TYPE field in Ethernet; it specifies which upper level protocol is in use. The most common PID now in use is F0 hex, which means "no upper level protocol, send data to the terminal." Much "packeting" goes on with no more protocol mechanism than this.

Higher level protocols

Lately there has been much experimentation with higher level protocols on top of AX.25. In addition to the DARPA Internet suite, several other efforts are underway. The most popular implementation at present is a commercial product of Software 2000, Inc., called NET/ROM. [Busch] A special ROM replaces the standard firmware in a TNC, turning it into a terminal switch. Users with conventional TNC software connect to a NET/ROM server node and may issue commands to connect to remote nodes and users. Internally, NET/ROM uses proprietary protocols: a connection-oriented "transport" protocol atop a connectionless network layer protocol. Automatic routing broadcasts establish network connectivity. Externally, however, NET/ROM presents only a connection-oriented service, concatenating its transport connections with a single AX.25 connection from each user. Compatibility with existing packet stations and software is maintained at the cost of sacrificing true end-to-end reliability and requiring the user to learn NET/ROM commands and some network topology information.

Other efforts

Other projects include "ROSE," [Beattie] an implementation of the packet layer from X.25, and TEXNET [McDermott]. The first components of ROSE are in the initial debugging stages, and TEXNET is being deployed, primarily in Texas and nearby areas.

TCP/IP on Amateur Packet Radio

The first amateur radio TCP/IP experiments were done by Richard Bisbey, NG6Q, with the ISI Amateur Radio Club, and Dave Mills, W3HCF. (Dave is well known to the Internet community as the creator of the Fuzzball, Network Time protocols, and a whole dictionary of colorful networking terms and expressions). Both efforts adapted existing IP gateway hardware and software to the task. Their work inspired me to create an Internet software package specifically designed for amateur packet radio use, running on hardware readily available to the average ham, and freely available in source form for study and experimentation. [Karn85A].

The result of this successful effort, to which several volunteers have made significant contributions, is the "KA9Q Internet Protocol Package." [Karn87] This software has now gained considerable operational experience. A class-A network address (44) has been assigned to the Amateur Packet Radio (AMPR) subnetwork and a domain name, **ampr.org**, has just been registered. This reflects the international nature of amateur radio; several dozen countries now have assignments in this address block. In the discussion that follows, I will refer to this emerging TCP/IP-on-amateur-radio network as AMPRNET (Amateur Packet Radio Network).

Amateur IP addressing

One important characteristic of AMPRNET is that the normally sharp distinction between Internet subnetworks is blurred. In the standard Internet model, each subnetwork is a well defined entity, e.g., an Ethernet LAN or Arpanet Wide Area Network (WAN), and each is fully interconnected internally, logically if not physically. In contrast, AMPRNET consists of individual stations and ad-hoc links, possibly unidirectional. Thus AMPRNET is a directed graph with individual stations (not fully connected subnetworks) at the nodes, and this required significant changes in IP address interpretation.

I call the approach I took "generalized subnetting" [Karn85B] since it goes beyond Internet standard subnetting, which only increased the number of allowable levels in the Internet addressing hierarchy from two to three. With generalized subnetting, the limit is 32, set by the number of bits in an IP address. This scheme was invented independently for the Fuzzball IP routers, and Comer calls it "Subnet Routing" [Comer, p 202]. The idea is that each entry in the routing table contains its own subnet mask. A routing lookup returns the matching table entry having the widest subnet mask.

There is *no* hardwired notion of an IP address "class" (A, B or C) and the subnet mask corresponding to a given address's entry may be different from one router to the next. This is a powerful technique, in that it allows the construction of an arbitrary network topology while still permitting the most compact routing table representation permitted by the address assignment strategy.

The ad-hoc "host specific" and "default" routing entries supported by most gateways are no longer necessary, since they are now merely special cases of a more general scheme. A host specific routing entry has a subnet mask of all ones, while the default route has a mask of all zeros. The default entry therefore matches any address, assuming some other entry (with wider mask) doesn't also match it.

Amateur Packet Radio and TCP/IP (*continued*)

The KA9Q Internet Protocol Package

Most Internet software packages for the PC generally support only one application (server or client) at a time. The ease of implementing custom applications is usually more important than the ability to support multiple simultaneous sessions, since other, larger systems running "real" operating systems (e.g., BSD UNIX and its derivatives) are usually available elsewhere on the net.

Suns and VAXen are still rare in the amateur radio world, however, so my goal was to create a usable network solely out of PC-class machines that could support at least the basic Internet applications simultaneously.

Because PCs running MS-DOS don't support multitasking, however, I chose to build a rudimentary form of multitasking into a single executable program containing the protocol modules. Unfortunately, this does make for an unusual programming environment. Each application is structured as a state machine driven by external events. For example, a TCP application is given three upcalls: receive data arrival, transmit data acknowledgement, and TCP state change. The application is responsible for maintaining its own state between upcalls, and it must process upcalls without hogging or blocking the processor.

Protocols

The KA9Q package supports the major DARPA Internet protocols: IP and ICMP at the Internet layer, TCP and UDP at the Host-Host (Transport) layer, and Telnet, FTP and SMTP at the Application layer. Multiple simultaneous server and client sessions are supported, with the only limit being available memory for control blocks and buffering. The package supports several subnet protocols and interfaces, with more under development. Ethernet (3Com's 3C501 interface) and SLIP (using PC asynchronous ports) are supported.

Since the package was developed specifically for AMPRNET, the AX.25 link level protocol is also incorporated. AX.25 may be used alone (for compatibility with amateur stations that don't yet support the Internet protocols) or as a subnet mechanism for carrying IP datagrams. The subnet interface is straightforward, with ARP [RFC 826] used to map IP addresses into AX.25 callsigns. ARP works automatically only when no digipeaters are involved; otherwise the mapping table must be set manually.

Encapsulation

AX.25 provides two forms of data encapsulation, the I (Information) and the UI (Unnumbered Information) frames. The former mode includes link level flow control and acknowledgement while the latter does not. This recognizes that the quality of packet radio channels that make up AMPRNET varies widely, ranging from fading, noisy 300 bps channels on long haul HF bands to 56 Kbps line-of-sight paths on UHF frequencies. Link level reliability mechanisms that are considered essential for performance on HF introduce intolerable overhead on reliable high speed links.

Type Of Service

The Type-of-Service (TOS) bits in each IP datagram select one mode or the other. (If no bits are set, a default mode is used). When I-frames are selected, the capability of fragmenting large datagrams at the link layer and immediately reassembling them at the other end of the link is also available.

Contributions to the Internet community

This follows the recommendation in [Kent] to avoid excessive IP level fragmentation by providing transparent subnet fragmentation whenever the subnet maximum transmission unit (MTU) would otherwise be unusually small. This facility is not available in UI mode, but this is acceptable since UI mode would only be used on high quality channels where large packets are already in use.

1200 baud amateur packet radio clearly represents the low end of the performance spectrum of subnetworks that now support the Internet protocols. In the words of Mike O'Dell (also a ham, with callsign N4NLN), AMPRNET currently holds TCP/IP's "worst wire" award. Frustration with the poor performance of existing TCP implementations over packet radio (primarily the fault of poor retransmission algorithms) motivated me to work in this area with Van Jacobson and Craig Partridge.

My experiences have made me somewhat of a radical, particularly when I see other implementors making arbitrary assumptions about the longest possible Internet round-trip time, or minimum subnet throughput, or maximum packet loss rate. The KA9Q TCP has the Van Jacobson congestion control algorithms, plus my own heuristic for ensuring the accuracy of round trip time measurements in the face of retransmissions, and they work well over the air. In keeping with the principle of "no arbitrary limits," there is no "give-up timer" or clipping of round trip time estimates. Retransmissions may occur indefinitely, subject to a back-off algorithm, of course; the decision to abort a connection is left up to the application or the human user.

Ad hoc internet

Because of its geographic dispersion, AMPRNET at present consists largely of isolated areas of activity. It is *not* interconnected with the DARPA Internet, nor does it use any other government resources, so low cost commercial facilities (e.g., Telenet's *PC Pursuit*) have instead been used to link these "islands" on a part time basis with good results. There seems to be considerable potential here to build an ad-hoc Internet out of available facilities, thus satisfying a major reason that amateur radio exists: public service, particularly in emergency communications.

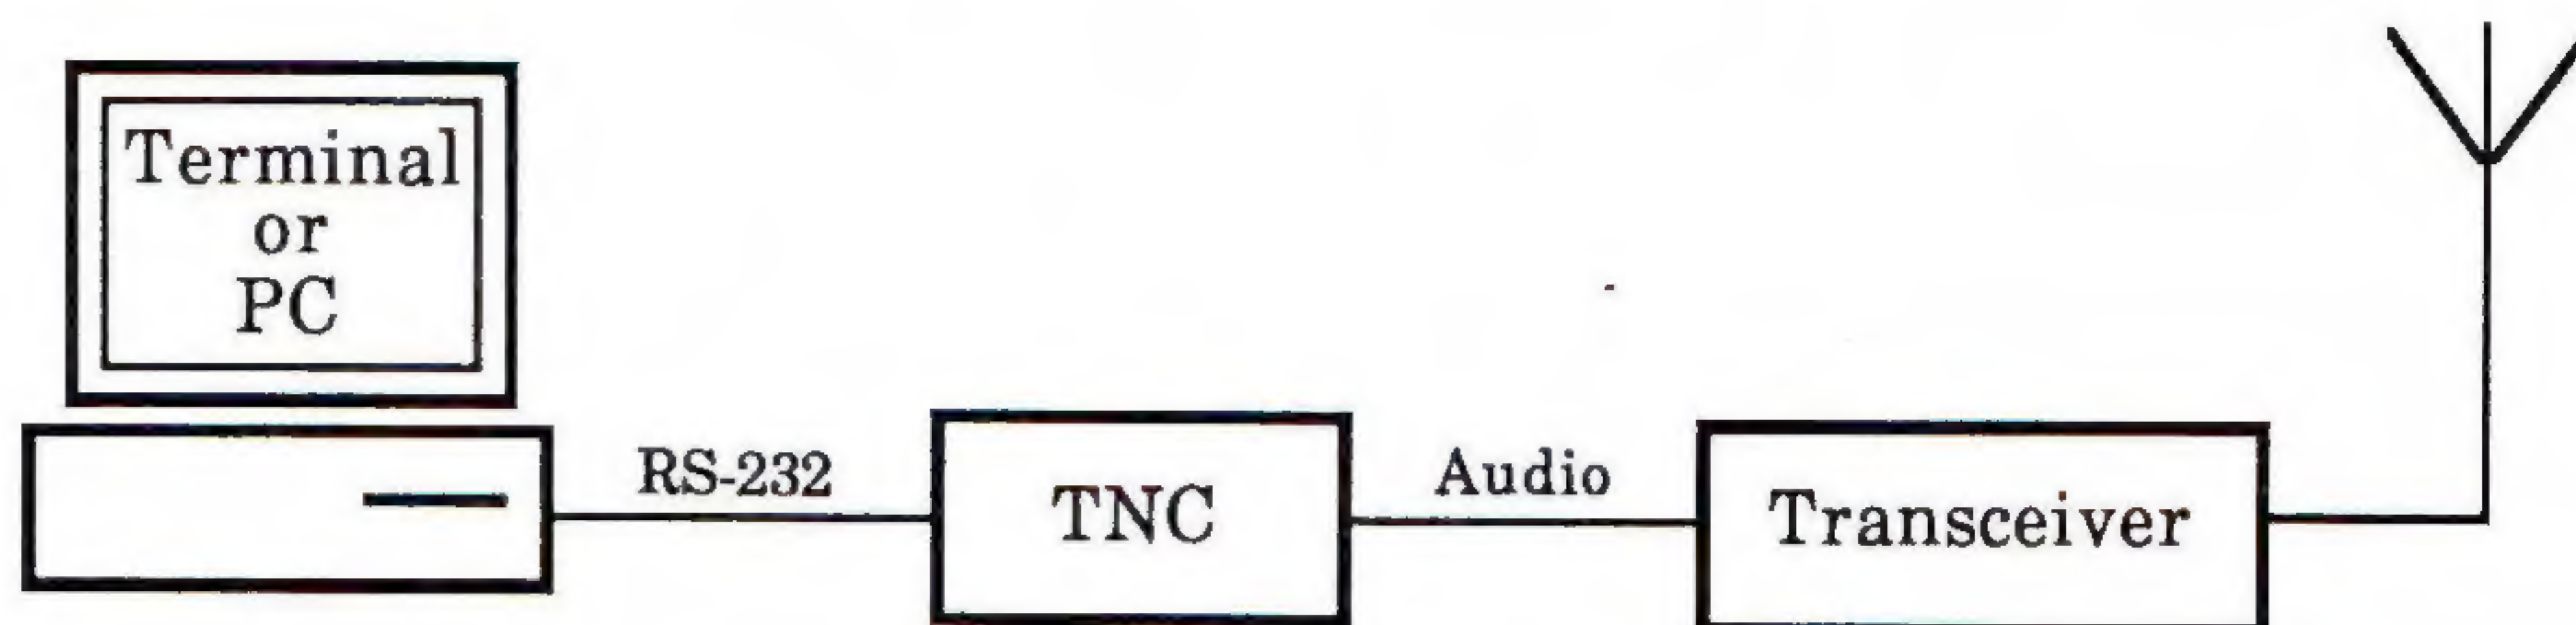
Our experiences with TCP/IP over amateur packet radio have contributed to the Internet effort by expanding the performance range of networks over which the protocols can operate. We have encountered and attacked problems that the rest of the Internet community has not had to face, but these solutions are now available should they be needed. In this way, amateur radio's charter for contributing to the state of the communication art is also fulfilled.

Availability of the KA9Q Internet Protocol Package

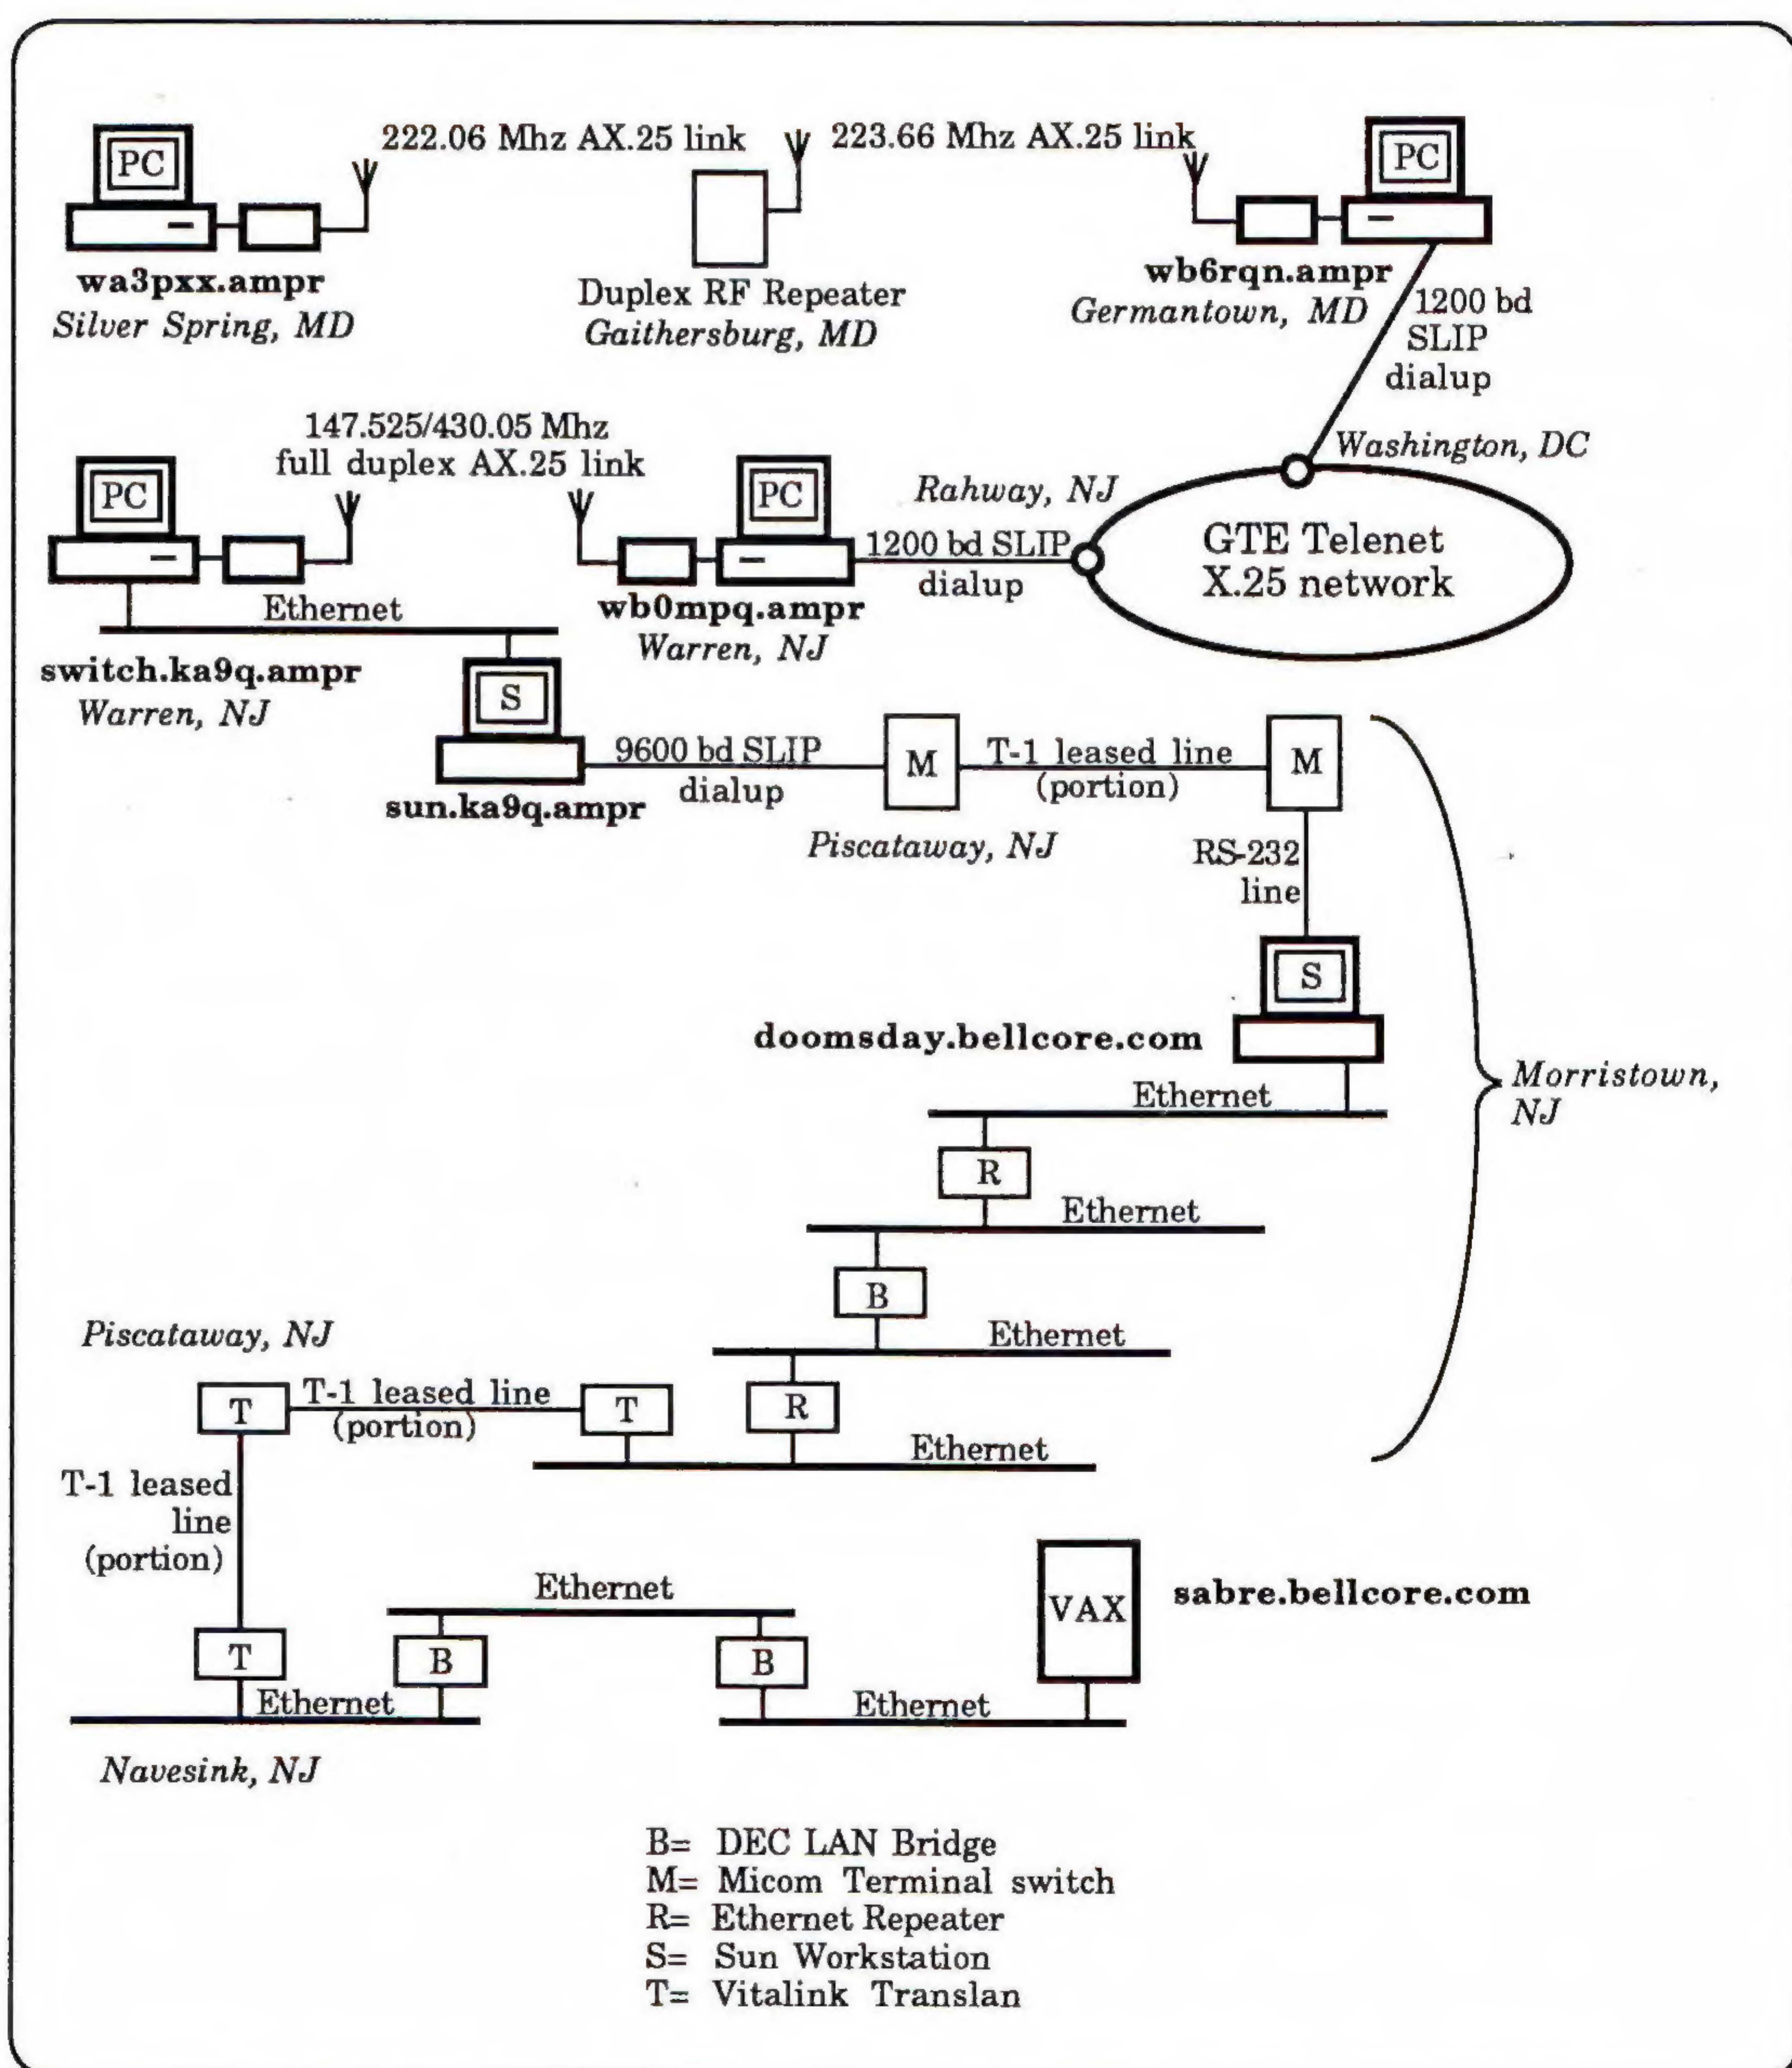
The KA9Q Internet Package is not public domain, but it is available by anonymous FTP for free noncommercial use. It may be obtained from [louie.udel.edu](ftp://louie.udel.edu) (10.0.0.96) under `/pub/ka9q` as a collection of ARC format archives. Source, executables and documentation is provided. Although the primary execution environment is MS-DOS on the IBM PC and clones, the code has been ported to a variety of machines including the Apple Macintosh, the Commodore Amiga, the Atari ST, and various flavors of UNIX System V. Although feedback and suggestions are welcome, this is a volunteer project so no warranties or guarantees can be offered.

continued on next page

Amateur Packet Radio and TCP/IP (continued)



Typical Amateur Packet Radio Station



This amateur "ad hoc internet" experiment was successfully conducted in January 1988.

References

- [ARRL82] Fox, T., ed., "Amateur Packet Radio Link Layer Protocol AX.25," American Radio Relay League.
- [Beattie] Beattie, G., "OSI: A Plan Comes Together," Proc. 6th ARRL Amateur Radio Computer Networking Conference. (ARCNC), 1987.
- [Busch] Busch, M. and Raikes, R., "NET/ROM Version 1.0 User's Manual," Software 2000, Inc. 1987.
- [Comer] Comer, D. "Internetworking with TCP/IP - Principles, Protocols, and Architecture," Prentice Hall, ISBN 0-13-470154-2.
- [Heatherington] Heatherington, D., "A 56 Kilobaud RF Modem," Proc. 6th ARRL ARCNC, 1987.
- [Karn] Karn, P., Price, H., and Diersing, R., "Packet radio in the Amateur Service," IEEE Journal on Selected Areas in Communications, May 1985.
- [Karn85A] Karn, P., "TCP/IP: A Proposal for Amateur Packet Radio Levels 4 and 3." Proc. 4th ARRL ARCNC, 1985.
- [Karn85B] Karn, P., "Addressing and Routing Issues in Amateur Packet Radio," Proc. 4th ARRL ARCNC, 1985.
- [Karn87] Karn, P., "The KA9Q Internet (TCP/IP) Package: A Progress Report." Proc. 6th ARRL ARCNC, 1987.
- [Kent] Kent, C., and Mogul, J., "Fragmentation Considered Harmful," Proc. ACM SIGCOMM 1987.
- [Levy] Levy, S., "Hackers - Heroes of the Computer Revolution," Anchor Press/Doubleday, 1984, ISBN 0-385-19195-2.
- [McDermott] McDermott, T., "Overview of the TEXNET datagram protocol," Proc. 6th ARRL ARCNC, 1987.
- [McGwier] McGwier, R., "DSP Modems: It's Only Software," Proc. 6th ARRL ARCNC, 1987.
- [RFC 826] Plummer, D., "An Ethernet Address Resolution Protocol."

PHIL KARN received his BSEE from Cornell University in 1978 and his MSEE (Computer Engineering) from CMU in 1979. After graduating, Phil worked at Bell Laboratories in Naperville, IL and Murray Hill, NJ. With the breakup of the Bell System in 1984, Phil went to work for Bellcore in Morristown, NJ. He is now a Member of the Technical Staff doing research into protocols for high speed fiber-based packet switching networks. Phil's interest in TCP/IP comes from two directions: amateur packet radio, and the large internet he helped build and manage within Bellcore. He has been a licensed radio amateur since 1971, and his interests there lie in the area of technical experimentation and development. Phil is on the board of directors of Tucson Amateur Packet Radio, the non-profit corporation largely responsible for making packet technology widely available to radio amateurs. He also sits on the Digital Communications Committee of the American Radio Relay League (ARRL), the US national amateur radio organization. In addition to terrestrial packet radio, Phil participates in the Radio Amateur Satellite Corporation (AMSAT), a group of volunteers that designs, builds and operates its own earth-orbiting satellites, including ones that feature packet switched communications.

CONNEXIONS

480 San Antonio Road
Suite 100
Mountain View, CA 94040

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

CONNEXIONS

PUBLISHER Daniel C. Lynch

EDITOR Ole J. Jacobsen

EDITORIAL ADVISORY BOARD Dr. Vinton G. Cerf, Vice President, National Research Initiatives.

Dr. David D. Clark, The Internet Architect, Massachusetts Institute of Technology.

Dr. David L. Mills, NSFnet Technical Advisor; Professor, University of Delaware.

Dr. Jonathan B. Postel, Assistant Internet Architect, Internet Activities Board; Division Director, University of Southern California Information Sciences Institute.

Subscribe to CONNEXIONS

U.S./Canada \$100. for 12 issues/year \$180. for 24 issues/two years \$240. for 36 issues/three years
International \$ 50. additional per year (Please apply to all of the above.)

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

☐ Check enclosed (in U.S. dollars made payable to CONNEXIONS).

☐ Charge my ☐ Visa ☐ Master Card Card # _____ Exp. Date _____

Signature _____

Please return this application with payment to:

CONNEXIONS

480 San Antonio Road Suite 100
Mountain View, CA 94040
415-941-3399

Back issues available upon request \$10./each
Volume discounts available upon request

CONNEXIONS